

Remarks

Entry of the amendments, reconsideration of the application, as amended, and allowance of all pending claims are respectfully requested. Claims 1-3 remain pending.

In the Office Action dated September 8, 2005, claims 1-3 are rejected under 35 U.S.C. 103(a) as being unpatentable over Sethuram et al. (U.S. Patent No. 5,828,903 in view of Ayanoglu et al. (U.S. Patent No. 6,122,759) in further view of Gentry et al. (U.S. Patent No. 5,778,180). Applicants respectfully, but most strenuously, traverse this rejection for the reasons below.

As indicated clearly in Applicants' specification, it is a feature of the claimed invention that not only is data transfer accomplished by what is conventionally referred to as a zero copy transfer, but that the claimed invention achieves this without the use of intermediate host buffers (specification page 4, line 25 through page 5 line 2). This is accomplished by transfer of the data directly into the address space of an end user, that is, someone at the application level. It is a well understood aspect of data transfer (I/O) in the computer arts that such operations are under the direct and continued control of operating system functions. In Applicants' invention some of this control is bypassed in a structured manner to achieve an even better zero copy transfer operation. A close reading of Applicants' specification by the Examiner will show that a significant portion of Applicants' description is taken up with describing the interface architecture structure that is used to achieve what one might call "zero copy transfer plus" or "**super** zero copy transfer." See Applicants' specification pages 10 through 15. The Examiner will find therein the detailed description of the interface mechanisms which permit application level users to establish the system states needed to carry out this function. Nothing akin to this interface is taught, disclosed or suggested by the cited art. Furthermore, the results achieved have been clearly appreciated as a significant advancement in the data transfer arts. While the art cited by the Examiner has similar goals, the cited art is nonetheless devoid of any mention of transferring data directly into any address space other than that belonging to the operating system. In contrast, Applicants have provided a structured mechanism for transferring data directly into the address space of a user at the application level.

It is clear from the Examiner's comments that there still remains a disagreement with respect to the meaning and import of the patent to Sethuram et al. It is Applicants' intent herein to attempt to explain the fact that there is an underlying fundamental difference between the teachings of this patent and that which is claimed as the present invention. The claims are modified herein to more clearly delineate this fundamental difference.

For emphasis, brevity and a clearer understanding, it is asserted that what is claimed is a true zero copy transfer, as that term is described in Applicants' specification. Such transfers avoid the use of intermediate buffers. Like Sethuram et al. the transfer is by DMA transfer to a host memory; however, unlike Sethuram et al. the transfer is to a "final" destination, that is, to a user's address space, not to a buffer from which the operating system or other software is forced to extract the data.

The Examiner has cited the following passage from the cited patent as a basis for asserting that Sethuram et al. teach Applicants' first claim step. This passage is reproduced below from the data found on the Patent and Trademark Office web site since it contains language which, when read properly, clearly illustrates the differences that the Applicants maintain as being most relevant to patentability.

In accordance with the teachings of the present invention, the adapter 202 controls the transfer of data between the host device 200 and the network 201 such that incoming cells are assembled directly in the host memory 208. Each incoming cell includes a header, which includes a field that identifies the virtual circuit associated with the given cell, and data, which contains the information intended for transmittal. For each incoming cell, the adapter 202 reads the header information and determines the appropriate buffer in the host memory 208 to send the cell data. Thus the incoming data is assembled directly in the host memory 208. [Emphasis added herein.]

It is critical to an understanding of the differences involved that the Examiner pay particular attention to the use of the word "cell" and all that the use of this term implies. A cell, as that term is defined and used by Sethuram et al. (col. 1, lines 29-31 and col. 4, lines 26-28, for example), is more than the data (message) that is intended for host system memory storage. Very importantly, a cell includes five bytes of header information indicating, among other things, the virtual register with which that cell is associated. This information is needed for subsequent processing of the data even though it has been transferred to the host

system memory. If the data portion were to be delivered to an ultimate destination in a user's address space, there would be no need whatsoever to include the five bytes of information that is appended to the data by Sethuram et al. to constitute what they call a "cell." It is clear from the passage above that cell data is transferred to the memory buffer. The cell data transferred includes this extra header information, which is not transferred in the claims of the present Applicants. Furthermore, as used by Sethuram et al., a buffer is a staging area for incoming data; it is not the final destination for that data.

The fact that Sethuram et al. contemplate further processing is also shown in the following passage from col. 1, lines 33-36 of the cited patent:

A receiver at the host device assembles these cells separately depending on which virtual circuit the incoming cell belongs to.

Clearly, the import of this passage is that Sethuram et al. contemplate that further processing is required subsequent to delivery of the "cells" to the host memory. **It is this very step that is avoided in the presently claimed invention.** To whatever extent that Sethuram et al. teach the use of real address information, it is the real addresses associated with intermediate buffers which the "host receiver" is capable of accessing for the purpose of data packet reassembly based on the header information which is contained in the incoming cells and which is transferred to the host memory.

Those of ordinary skill in the art would not be led to a process which includes a step which is necessary in the method of Sethuram et al. to accomplish what Applicants have done without this step. The fundamental patent upon which the Examiner relies to support the rejection is actually thus seen to teach away from that which Applicants have claimed. The other two patents cited in support of the rejection cannot and do not compensate for this fundamental difference.

With respect to the patent to Ayanoglu et al. it is noted that, while they appear to teach the use of an acknowledgement step, they fail to teach anything with respect to the transfer of data directly into an application level address space. This is a feature that is also lacking in the patent to Sethuram et al., as pointed out above. Ayanoglu et al. are focused on the ATM (Asynchronous Transfer Method) protocol. Nowhere in that protocol is there any

reference to the transfer of data directly into an application level address space. Furthermore, it is noted that there is no mention whatsoever in the patent to Ayanoglu et al. of the term “address space” or “direct memory access (DMA).

With respect to the patent to Gentry et al. the following passage is relevant (col. 2, lines 46-63):

More specifically, each ATM connection may have a private pool of buffers, into which only packets for that connection will be placed. Since the pool of buffers is private, a program can be given access to its own pool. No data copying will be required for packets received into the private pool. Therefore, a packet may be directly sent to its final destination by DMA. Additionally, protected buffer descriptors prevent corruption of data with the private buffers dedicated to the data's final destination. When a packet arrives, if there are no private buffers available, the router falls back to a common pool of buffers which are not available to the programs and thus, must be copied. Since not all connections will be able to use private buffer pools due to lack of resources, a change in the connection from the common pool of buffers to the private pool of buffers and vice versa is available. This change affects a connection while it operates. The change takes effect on the next packet to arrive. [Emphasis added herein.]

It is noted that, while it appears that the goals of Gentry et al. are similar to those found in the present invention, the methods employed to achieve them are different. Again, there is no writing of data directly into a user's address space. Instead, application programs are given access to separate “private buffers.” While this is a laudable goal with similar objectives, it is not the same process. There is no suggestion that these private buffers are within the user's address space where transferred data can be used immediately without special arrangements to assure access to data within these private buffers. As is clearly stated in the passage above, programs can be given access to these private buffers. In stark contrast, in the claimed invention user level applications (programs) automatically have access to the incoming data since it is placed directly into the user's address space.

It is also noted from the passage from Gentry et al. above that they contemplate the situation in which private buffers are not available. If the data were to be written directly into a user's address space such a situation would not arise. It thus becomes clear that Gentry

et al., while having similar goals, fail to teach disclose or suggest the writing of data directly into a user's address space, as is specifically recited in Applicants' claims 1-3.

It is therefore seen from the above that the indicated rejection is not supported by a proper reading of the teachings from Sethuram et al. and/or the other two cited patents. Accordingly, it is respectfully requested that the rejection of Applicants' claims 1-3 based on the patent to Sethuram et al. and the other cited patents be withdrawn.

As a concluding matter, it is noted that support for the changes made to Applicants' claims can be found throughout the specification. However, for clarity, it is noted that on page 6, starting on line 15, the following description is found which fully supports the subject claim modifications:

It is still a further object of the present invention to enable users running applications in their own address spaces on one data processing system to be able to transfer data accurately and efficiently into a user's address space in another data processing system whether or not that data processing system is remote or in fact contained within the same physical package or frame.

Should the Examiner wish to discuss this case with applicants' attorney, please contact applicants' attorney at the below listed number.

Respectfully submitted,



Lawrence D. Cutter
Attorney for Applicants
Registration No.: 28,501

Dated: December 8, 2005.

HESLIN ROTHENBERG FARLEY & MESITI P.C.
5 Columbia Circle
Albany, New York 12203-5160
Telephone: (518) 452-5600
Facsimile: (518) 452-5579